arXiv:1512.06195v1 [cs.DL] 19 Dec 2015

Quantifying Orphaned Annotations in Hypothes.is

Mohamed Aturban, Michael L. Nelson, and Michele C. Weigle

Dept of Computer Science, Old Dominion University, Norfolk, VA, 23529
{maturban,mln,mweigle}@cs.odu.edu
http://ws-dl.cs.odu.edu

Abstract. Web annotation has been receiving increased attention recently with the organization of the Open Annotation Collaboration and new tools for open annotation, such as Hypothes.is. We investigate the prevalence of *orphaned annotations*, where neither the live Web page nor an archived copy of the Web page contains the text that had previously been annotated in the Hypothes.is annotation system (containing 20,953 highlighted text annotations). We found that about 22% of highlighted text annotations can no longer be attached to their live Web pages. Unfortunately, only about 12% of these annotations can be reattached using the holdings of current public web archives, leaving the remaining 88% of these annotations orphaned. For those annotations that are still attached, 53% are in danger of becoming orphans if the live Web page changes. This points to the need for archiving the target of an annotation at the time the annotation is created.

Keywords: Web Annotation, Web Archiving, HTTP

1 Introduction

Annotating web resources helps users share, discuss, and review information and exchange thoughts. Haslhofer et al. [12] define annotation as associating extra pieces of information with existing web resources while the Open Annotation Collaboration (OAC) group defines an annotation as a set of connected resources [12] where the basic form of annotation consists of the Body and Target resources. Ideally, the Body should be about the Target. Annotation types include commenting on a web resource, highlighting text, replying to others' annotations, specifying a segment of interest rather than referring to the whole resource, tagging, etc.

Hypothes.is[1], an open annotation tool, was released in early 2013 and is publicly accessible for users to annotate, discuss, and share information. It provides different ways to annotate a web resource: highlighting text, adding notes, and commenting on and tagging a web page. In addition, it also allows users to share

This arXiv paper is an extended version of our TPDL 2015 paper [3].

[1] http://hypothes.is

an individual annotation URI with each other as an independent web resource. The annotation is provided in JSON format and includes the annotation author, creation date, target URI, annotation text, permissions, tags, comments, etc.

One of the well-known issues of the Web is that Web pages are not fixed resources. A year after publication, about 11% of content shared on social media will be gone [15, 16], and 20% of scholarly articles have some form of reference rot [13]. Lost or modified web pages may result in *orphaned annotations*, which can no longer be attached to their target web pages.

Figure 1 shows the annotated web page `http://caseyboyle.net/3860/readings/against.html` which has 144 annotations from Hypothes.is. The text with darker highlights indicates more users have selected this part of the page to annotate. The issue here is that all of these annotations are in danger of being orphaned because no copies of the target URI are available in the archives. Figure 2 shows the target URI `http://climatefeedback.org/`, created in December 2014, with the annotation "After reading about your project at MIT news, I visited your page and ..." on the highlighted text "Scientific feedback for Climate Change information online". In August 2015, this annotation can no longer be attached to the target web page because the highlighted text no longer appears on the page, as shown in Figure 3. Although the live Web version of `http://climatefeedback.org/` has changed and the annotation was in danger of being orphaned, the original version that was annotated has been archived and is available at the Internet Archive (`https://web.archive.org/web/20141210121018/http://climatefeedback.org/`). The annotation could be re-attached to this archived resource.

This paper is a follow-up to our paper [3] presented in August 2015 in the 19th International Conference on Theory and Practice of Digital Libraries (TPDL). In the TPDL version, we analyzed 6281 highlighted text annotations collected in January 2015, while in this paper, we worked with 20,953 annotations collected in August 2015. Figure 4 shows that the number of annotations in Hypothes.is has increased since July 2013. In this paper, we introduce a detailed analysis of the extent of orphaned highlighted text annotations in the Hypothes.is annotation system as of August 2015. We also look at the potential for web archives to be used to reattach these annotations. We find that 22% of the highlighted text annotations at Hypothes.is are not attached to the live web, and only a few can be reattached using web archives. Further, we show that 53% of the currently attached annotations could potentially become orphans if their live Web resources change, because there are no archived versions of the annotated resources available. Our analysis points to the potential for reducing orphaned annotations by archiving web resources at the time of annotation.

2 Related Work

Annotation has long been recognized as an important and fundamental aspect of hypertext systems [14] and an integral part of digital libraries [1], but broad adoption of general annotation for the Web has been slow. Annotations have been

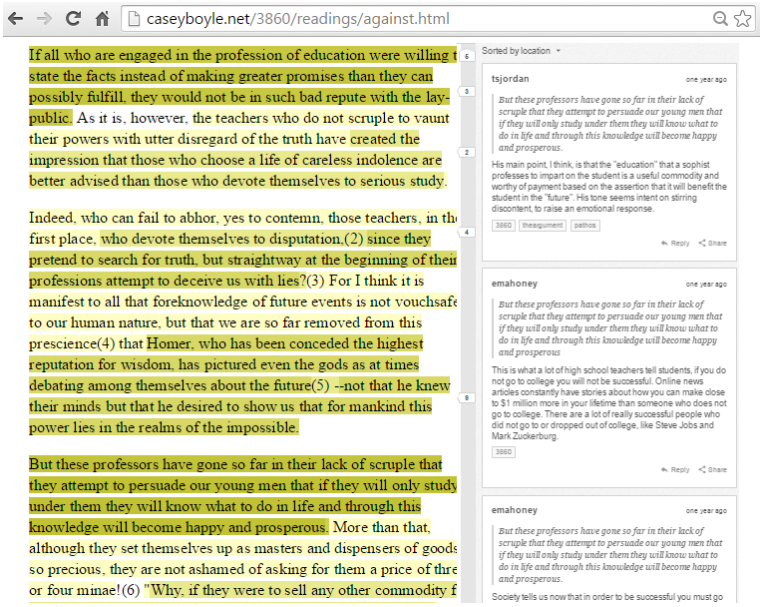
caseyboyle.net/3860/readings/against.html
Sorted by location
tsjordan
one year ago
But these professors have gone so far in their lack of scruple that they attempt to persuade our young men that if they will only study under them they will know what to do in life and through this knowledge will become happy and prosperous.
His main point, I think, is that the "education" that a sophist professes to impart on the student is a useful commodity and worthy of payment based on the assertion that it will benefit the student in the "future". His tone seems intent on stirring discontent, to raise an emotional response.
3860
theargument
pathos
Reply
Share
emahoney
one year ago
But these professors have gone so far in their lack of scruple that they attempt to persuade our young men that if they will only study under them they will know what to do in life and through this knowledge will become happy and prosperous
This is what a lot of high school teachers tell students, if you do not go to college you will not be successful. Online news articles constantly have stories about how you can make close to $1 million more in your lifetime than someone who does not go to college. There are a lot of really successful people who did not go to or dropped out of college, like Steve Jobs and Mark Zuckerburg.
3860
Reply
Share
emahoney
one year ago
But these professors have gone so far in their lack of scruple that they attempt to persuade our young men that if they will only study under them they will know what to do in life and through this knowledge will become happy and prosperous.
Society tells us now that in order to be successful you must go
If all who are engaged in the profession of education were willing t
state the facts instead of making greater promises than they can
possibly fulfill, they would not be in such bad repute with the lay-
public. As it is, however, the teachers who do not scruple to vaunt
their powers with utter disregard of the truth have created the
impression that those who choose a life of careless indolence are
better advised than those who devote themselves to serious study.
Indeed, who can fail to abhor, yes to contemn, those teachers, in th
first place, who devote themselves to disputation,(2) since they
pretend to search for truth, but straightway at the beginning of their
professions attempt to deceive us with lies?(3) For I think it is
manifest to all that foreknowledge of future events is not vouchsafe
to our human nature, but that we are so far removed from this
prescience(4) that Homer, who has been conceded the highest
reputation for wisdom, has pictured even the gods as at times
debating among themselves about the future(5) --not that he knew
their minds but that he desired to show us that for mankind this
power lies in the realms of the impossible.
But these professors have gone so far in their lack of scruple that
they attempt to persuade our young men that if they will only study
under them they will know what to do in life and through this
knowledge will become happy and prosperous. More than that,
although they set themselves up as masters and dispensers of goods
so precious, they are not ashamed of asking for them a price of thre
or four minae!(6) "Why, if they were to sell any other commodity f

Fig. 1: Using the Hypothes.is Browser Extension to View the Annotations of http://caseyboyle.net/3860/readings/against.html

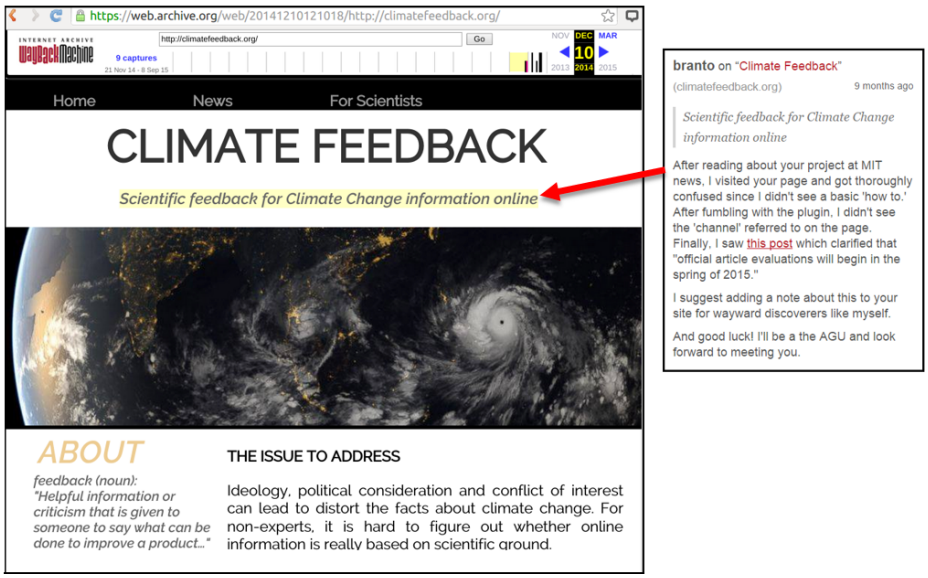
https://web.archive.org/web/20141210121018/http://climatefeedback.org/
http://climatefeedback.org/
Go
9 captures
Home
News
For Scientists
CLIMATE FEEDBACK
Scientific feedback for Climate Change information online
ABOUT
feedback (noun):
"Helpful information or criticism that is given to someone to say what can be done to improve a product..."
THE ISSUE TO ADDRESS
Ideology, political consideration and conflict of interest can lead to distort the facts about climate change. For non-experts, it is hard to figure out whether online information is really based on scientific ground.
branto on "Climate Feedback"
(climatefeedback.org)
9 months ago
Scientific feedback for Climate Change information online
After reading about your project at MIT news, I visited your page and got thoroughly confused since I didn't see a basic 'how to.' After fumbling with the plugin, I didn't see the 'channel' referred to on the page. Finally, I saw this post which clarified that "official article evaluations will begin in the spring of 2015."
I suggest adding a note about this to your site for wayward discoverers like myself.
And good luck! I'll be a the AGU and look forward to meeting you.

Fig. 2: http://climatefeedback.org/ in December 2014

climatefeedback.org
Climate Feedback
A voice for science in climate change media coverage
Today's media climate leads to confusion
With so much information available online, trying to figure out which information is credible — and what is not — is a real challenge. When so much of what we read falls outside of our own expertise, how can we know which headlines and news articles are consistent with science?
theguardian
Climate change
Climate change should be top foreign policy
Forbes
New Posts
Most Popular
Lists
Larry Bell, Contributor
branto on "Climate Feedback"
(climatefeedback.org)
9 months ago
Scientific feedback for Climate Change information online
After reading about your project at MIT news, I visited your page and got thoroughly confused since I didn't see a basic 'how to.' After fumbling with the plugin, I didn't see the 'channel' referred to on the page. Finally, I saw this post which clarified that "official article evaluations will begin in the spring of 2015."
I suggest adding a note about this to your site for wayward discoverers like myself.
And good luck! I'll be a the AGU and look forward to meeting you.

Fig. 3: http://climatefeedback.org/ in August 2015

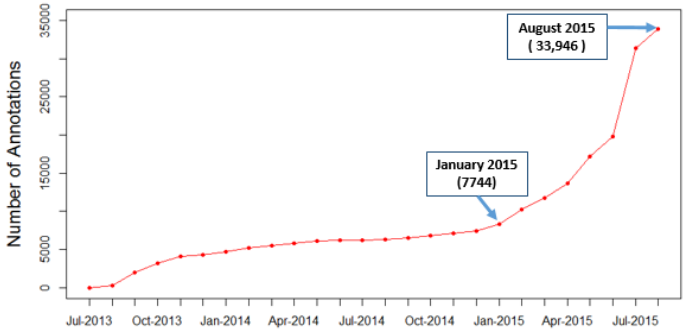

Fig. 4: Number of Annotations in Hypothes.Is since July 2013

studied for digital library performance [9, 20] and methods have been explored for aligning annotations in modified documents [7], but typically such studies are limited to annotation systems specific for a particular digital library. While orphaned annotations of general web pages have been studied in the context of Walden's Paths [10, 8], our study of Hypothes.is is a more recent evaluation of annotation and page synchronization in a widely deployed system.

Memento [22] is an HTTP protocol extension that aggregates information about the resources available in multiple Web archives. We can use Memento to obtain a list of archived versions of resources, or mementos, available in several web archives. In this paper, we use the following Memento terminology:

- URI-R - the original resource as it used to appear on the live Web. A URI-R may have 0 or more mementos (URI-Ms)
- URI-M - an archived snapshot of the URI-R at a specific date and time, which is called the Memento-Datetime, *e.g.*, $URI\text{-}M_i = URI\text{-}R@t_i$
- TimeMap - a resource that provides a list of mementos (URI-Ms) for a URI-R, ordered by their Memento-Datetimes

There has been previous work in developing annotation systems to support collaborative work among users and in integrating the Open Annotation Data Model [18] with the Memento framework. The Open Annotation Collaboration (OAC) [12] has been introduced to make annotations reusable through different systems like Hypothes.is. Before publishing OAC, annotations would not be useful if the annotated web pages were lost because annotations were not assigned URIs independent from the web pages' URIs. By considering annotations as first-class web resources with unique URIs, annotation not only would become reusable if their targets disappear, but also would support interactivity between systems. Sanderson and Van de Sompel [19] built annotation systems that support making web annotations persistent over time. They focus on integrating features in the Open Annotation Data Model with the Memento framework to help reconstructing annotations for a given memento and retrieving mementos for a given annotation. They did not focus on the case of orphaned annotations and assumed that the archived resources were available in web archives. Ainsworth et al. have estimated how much of the web is archived [2]. The result indicated that 35-90% of publicly accessible URIs have at least one archived copy, although they did not consider annotations in their work, the result might estimate the number of orphaned annotations by factors like how frequently web pages are archived and the archiving process coverage. In other work [17, 11] researchers built annotation systems that can deliver a better user experience for specialized users and scholars. The interfaces allow users to annotate multimedia web resources as well as medieval manuscripts in a collaborative way. In this paper, we focus on orphaned annotations and investigate how web archives could be used to reattach these annotations to the original text.

3 Methodology

We performed our analysis on the publicly accessible annotations available at Hypothes.is as of August 2015. The interface allows users to create different types of annotations: (1) making a note by highlighting text and then adding comments and tags about the selected text, (2) creating highlights only, (3) adding comments and tags without highlighting text, and (4) replying to an existing annotation. Table 1 shows how many annotations belong to each type.

In August 2015, we downloaded the JSON of all 33,946 publicly available annotations from Hypothes.is. Figure 5 shows the JSON of the annotation from Figure 2 with relevant fields shown in bold. The "`updated`" field gives the annotation creation date, "`source`" provides the annotation target URI, "`type`" : "`TextQuoteSelector`" indicates that it is a highlighted text annotation,

Number of Annotations	Highlighted Text	Notes	Tags
11,289		✓	
9858	✓	✓	
9252	✓	✓	✓
1835	✓		✓
1356		✓	✓
348			✓
8	✓		
Total (33,946)			

Table 1: Annotation Types in Hypothes.is

Number of Annotations	Host
1222	`caseyboyle.net`
1191	`www.perseus.tufts.edu`
887	`rhetoric.eserver.org`
875	`networkedlearningcollaborative.com`
749	`sosol.perseids.org`
733	`tkbr.ccsp.sfu.ca`
526	`shakespeare.mit.edu`
391	`hypothes.is`
356	`renaissancejohnson.weebly.com`
336	`moodle2.wesleyan.edu`

Table 2: The Top Hosts with Annotated Pages

"`exact`" contains the highlighted text, and "`text`" contains the annotation text itself. We focus only on annotations with highlighted text ("`type`":"`TextQuoteSelector`"), leaving 20,953 annotations for analysis. To determine how many of those annotations are orphaned, for each annotation we performed the following steps:

- Determine the current HTTP status of the annotation target URIs ("`source`").
- Compare selected highlighted text ("`exact`") to the text of the current version of the URI.
- Discover available mementos for the target URI.
- Search for highlighted text within the discovered mementos.

In Table 2, we show the top 10 hosts with annotations at Hypothes.is. Many of these hosts, including the top three, are academic servers and appear to use the system for annotation of scholarly work. Apart from this listing, we did not attempt to make judgements about the content of the annotations or annotation target text in our analysis.

```
{
  "updated": "2014-12-03T04:47:21.863568+00:00",
  "group": "__world__",
  "target": [
      {
          "scope": ["http://climatefeedback.org"],
          "selector": [{
                  "endContainer": "/div[2]/p[1]",
                  "endOffset": 57,
                  "type": "RangeSelector",
                  "startOffset": 0,
                  "startContainer": "/div[2]/p[1]"
              },{
                  "start": 50,
                  "end": 107,
                  "type": "TextPositionSelector"
              },{
                  "exact": "Scientific feedback for Climate
                            Changeinformation online",
                  "prefix": "For Scientists CLIMATE FEEDBACK",
                  "type": "TextQuoteSelector",
                  "suffix": "ABOUT feedback (noun): Ḧelpful"
              }
          ],
          "pos": {"top": 148,"height": 25},
          "source": "http://climatefeedback.org/"
      }
  ],
  "tags": [],
  "text": "After reading about your project at MIT news,
           I visited your page and got thoroughly confused
           since I didn't see a basic 'how to.' After
           fumbling with the plugin, I didn't see the
           'channel' referred to on the page.  Finally,
           I saw [this post] (http://oceans.mit.edu/
           featured-stories/ climate-feedback)which
           clarified that "official article evaluations
           will begin in the spring of 2015."I suggest
           adding a note about this to your site for
           wayward discoverers like myself.And good
           luck! I'll be a the AGU and look forward
           to meeting you. ",
  "created": "2014-12-03T04:46:57.630434+00:00",
  "uri": "http://climatefeedback.org/",
  "user": "acct:branto@hypothes.is",
  "document": { ... },
  "consumer": "00000000-0000-0000-0000-000000000000",
  "id": "xNec2gjYT5-ORcDg4fl7nA",
  "permissions": {
      "admin": ["acct:branto@hypothes.is"],
      "read": ["acct:branto@hypothes.is","group:__world__"],
      "update": ["acct:branto@hypothes.is"],
      "delete": ["acct:branto@hypothes.is"]
  }
}
```


Fig. 5: An Annotation Described in JSON Format, Available at https://hypothes.is/api/annotations/xNec2gjYT5-ORcDg4fl7nA

3.1 Determining the HTTP Status

In the first step, the current HTTP status of annotation target URIs can be obtained by issuing HTTP HEAD requests for all URIs. In addition, we extended

this to detect "soft" 401, 403, and 404 URIs, which return a 200 OK status but actually indicate that the page is not found or is located behind authentication [4]. One technique we used to detect "soft" 4xx is to modify the original URI by adding some random characters to the parent directory, so that it is likely that the new URI does not exist. After that, we download the content of the original URI and the new one. If the content of both web pages is 93% (or above) identical, and the HTTP status code of these URIs is "200 OK", then we consider that the HTTP status of the original URI is a "soft" 4xx. We have written a Python program available in GitHub[2] for detecting "soft" 4xx URIs.

The returned responses will determine the next action which should be made for every URI. The resulting responses can be categorized into three different groups. The first group contains URIs with hostnames `localhost` or URIs which are actually URNs. The second group has URIs with one of the following status codes: "soft" and actual 400, 401, 403, 404, 429 or Connection-Timeout. URIs with 200 status code belong to the third group.

The first group, localhost and URNs, were excluded completely from our analysis because these are pages that are not publicly accessible on the live Web. URIs in the second group, soft/actual 4xx and timed-out URIs, have been checked for mementos in the web archives. For URIs with response code 200, we have compared their associated highlighted annotation text with both the current version of the web page and the available mementos in the archives. Even though some annotations are still attached to their live web pages, we are interested to see if they have mementos to know how likely those annotations are to become orphans if their current web pages change or become unavailable.

3.2 Are Annotations Attached to the Live Web?

The second step is to compare the annotated text ("`exact`") of each annotation target URI that has a 200 HTTP status code with the current version of its web page and see if they match; this can be done by downloading the web page and extracting only the text which will be compared to the highlighted annotation text. Extracting text from a standard HTML resource is different from doing it with a Portable Document Format (PDF) file. We use `curl` to access and download web pages. Then, for HTML web resources, we extract only the text after cleaning it by removing all HTML tags, extra white-space characters, and others. If the web resource is a PDF file, extracting text will need one more step of converting the original binary contents to plain text. `pdftotext`[3], which runs from the command-line and is freely available in many Linux distributions, can extract plain text from PDF files. After extracting the text either from a standard HTML web page or a PDF file, we search for the highlighted annotation text. If the text is not found, the annotation is considered *not attached*. For example, as shown in Figure 3, the annotation text is no longer attached to the web page as the highlighted text, shown in Figure 2, has been removed from the live web page.

[2] `https://github.com/maturban/Soft_4xx`
[3] http://www.foolabs.com/xpdf/download.html

For all annotations that are not attached until this point, we use PhantomJS to download the web page again trying to capture more resources that are loaded by JavaScript [5, 6].

It is important to mention that when examining the existence of the annotated text in the related web pages, there are false negatives. In other words, our program might not identify some truly attached annotations because of reasons including but not limited to, not detecting "soft" 4xx URIs, failure in extracting text from PDF files by `pdftotext`, missing embedded resources in web pages that are loaded by JavaScript even with the use of PhantomJS. For example, the annotation that is about the highlighted equation in Figure 6 can not be re-attached to the the PDF file[4] because the highlighted equation is saved with a different encoding. Figure 7 illustrates another example where a web page[5] needs authentication to access the missing content. We can consider this example as an undetected "soft" 4xx as the server in this case should respond with 401 HTTP status code, but it returns "200 OK" instead.

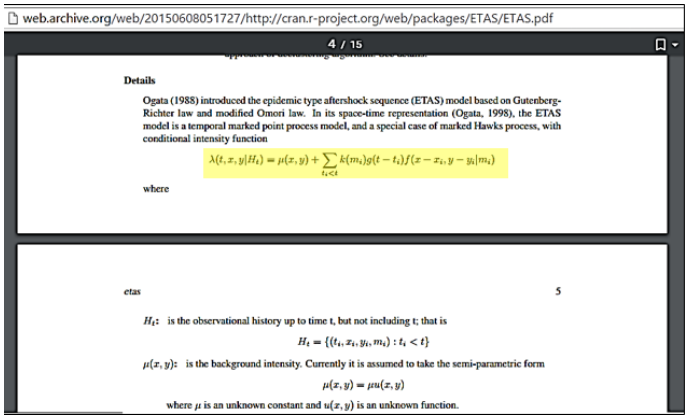
web.archive.org/web/20150608051727/http://cran.r-project.org/web/packages/ETAS/ETAS.pdf
4 / 15
Details
Ogata (1988) introduced the epidemic type aftershock sequence (ETAS) model based on Gutenberg-Richter law and modified Omori law. In its space-time representation (Ogata, 1998), the ETAS model is a temporal marked point process model, and a special case of marked Hawks process, with conditional intensity function
$\lambda(t,x,y|H_t) = \mu(x,y) + \sum_{t_i<t} k(m_i)g(t-t_i)f(x-x_i, y-y_i|m_i)$
where
etas
5
H_t: is the observational history up to time t, but not including t; that is
$H_t = \{(t_i, x_i, y_i, m_i) : t_i < t\}$
$\mu(x,y)$: is the background intensity. Currently it is assumed to take the semi-parametric form
$\mu(x,y) = \mu u(x,y)$
where μ is an unknown constant and $u(x,y)$ is an unknown function.

Fig. 6: Example of an Annotation That Can Not Be Re-attached to a PDF File (http://web.archive.org/web/20150608051727/http://cran.r-project.org/web/packages/etas/etas.pdf)

3.3 Discovering Mementos for All Valid URIs

The third step is to discover mementos for all valid annotation target URIs. For this purpose, we used a Memento Aggregator [21, 22] which provides a TimeGate through which we can get mementos that are closest to a URI-R's date. It would

[4] `http://web.archive.org/web/20150608051727/http://cran.r-project.org/web/packages/ETAS/ETAS.pdf`

[5] `http://proquest.safaribooksonline.com/book/programming/scala/9781449368814/preface/_who_this_book_is_for_html`

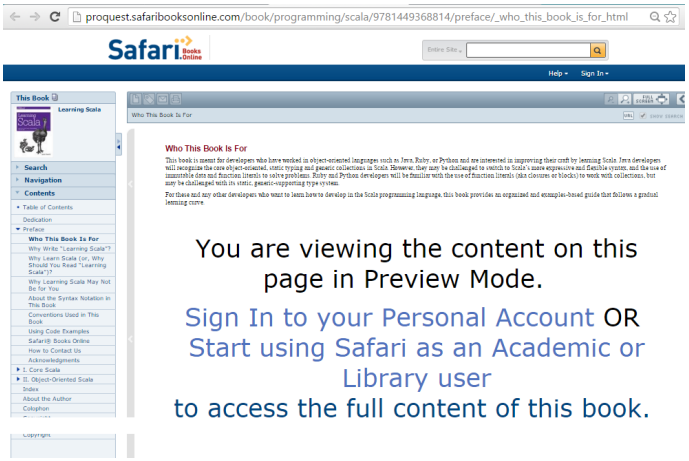
proquest.safaribooksonline.com/book/programming/scala/9781449368814/preface/_who_this_book_is_for_html
Safari
Who This Book Is For
You are viewing the content on this page in Preview Mode.
Sign In to your Personal Account OR
Start using Safari as an Academic or Library user
to access the full content of this book.

Fig. 7: Example of a Web Page (http://proquest.Safaribooksonline.Com/book/programming/scala/9781449368814/preface/_who_this_book_is_for_html) That Needs Authentication

be a time-consuming task to check all available mementos for a URI-R to see whether they can be used to recover web pages. For example, URIs like `http://www.nytimes.com/` or `http://www.cnn.com/` have thousands of mementos in different archives. The strategy that we use here is efficient in terms of execution time. For each URI, we only retrieve the nearest mementos to the annotation's creation date ("`updated`"). More precisely, we are capturing the closest memento to the date *before* the annotation was created and the closest memento to the date *after* the annotation was created.

In Figure 8(a), the annotation A was created at the time t_5. The closest memento to the date before t_5 was M_2 (captured at t_4) while the closest memento to the date after t_5 was M_3 (captured at t_7). So, for this annotation we picked the two closest mementos, which are M_2 and M_3. Figure 8(b) is an example where mementos are only available before the annotation creation date while in Figure 8(c), mementos are only available after the annotation creation date. It is also possible that an annotation target has no mementos at all, as Figure 8(d) shows. If there are multiple closest mementos from different archives that share the same creation date (*memento-datetime*), then we consider all of these mementos for two different reasons. First, it is possible that at the time a memento is requested from an archive, there would be a technical problem or server-related issue which may affect returning the requested mementos. Second, we would like to know how different archives could contribute to provide mementos and recover annotation target text.

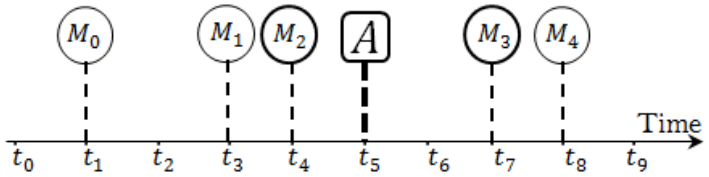

(a) Existing Mementos Before and After the Annotation Creation Date

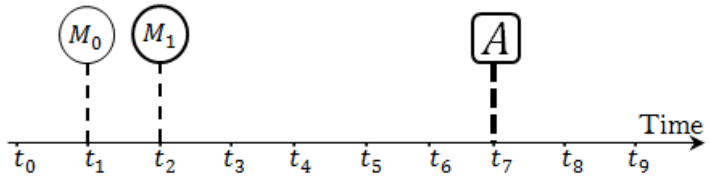

(b) Existing Mementos Only Before the Annotation Creation Date

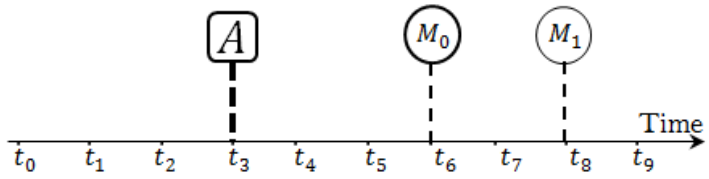

(c) Existing Mementos Only After the Annotation Creation Date

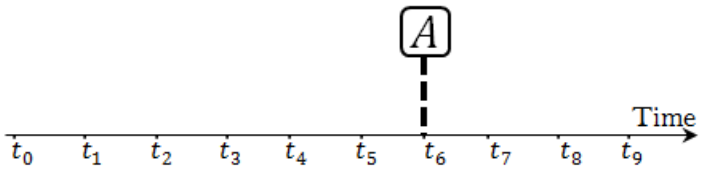

(d) No Mementos for the Annotation Target

Fig. 8: Annotation and Memento Creation Dates

3.4 Are Annotations Attached to the Selected Mementos?

The final step is to see whether annotated URIs can be recovered by their mementos. The same technique introduced in Section 3.2 is used to test mementos. If the annotation target text ("`exact`") matches the text in the discovered memento, then we consider that this annotation is attached to the memento. Otherwise, we consider that the annotation cannot be attached.

4 Results

We collected 33,946 annotations from Hypothes.is. Table 3 shows the results of checking the HTTP status code for the target URIs of all 20,953 highlighted text

Number of Annotations	Status Code
18167	200
778	Time out
666	404
326	`localhost`
318	URN
190	Soft 401/403/404
176	Unknown
87	401
80	503
68	403
48	410
21	406
19	500
5	416
2	520
1	400
1	504

Table 3: HTTP Status Code for All Annotation Target URIs

annotations. We noticed that a group of 820 annotations should be excluded from our analysis. This group consists of annotations whose target URIs are unresolvable, such as `localhosts` and URNs. In our further analysis, we will focus only on the 20,133 annotations that contain highlighted text and have resolvable target URIs. We noticed also that out of 20,133 annotations, 10% (1966) have URI-Rs that are no longer available on the live web, returning 400 and 500 status codes. Figure 9 shows how annotations are classified based on the status codes of their target URIs.

After checking each annotation, we found that 15,773 (78%) of the highlighted text annotations are still attached to their live web pages as Figure 10 shows. This means that the remaining 22% of the annotations are orphans if there are no mementos that can be used to reattach these annotations.

Next for each annotation, we checked the archives for the presence of mementos of the target URI near the annotation creation date. In Table 4 we consider annotations that have mementos both before ("L") and after ("R") the annotation date. "No" under the L and R columns means that annotation cannot be attached to the nearest memento, while "Yes" means that the annotation attaches to the nearest memento.

Table 5 shows the number of annotations that have mementos only on the L side (before) of the annotation date, and Table 6 shows the number of annotations that have mementos only on the R side (after) of the annotation date. Finally, Table 7 illustrates the number of annotations whose targets have no mementos. From these tables, we see that 4360 (22%) of the annotations can no longer be attached to their live web pages. Unfortunately, the current holdings of web archives only allow 12% of these to be re-attached, while the remaining

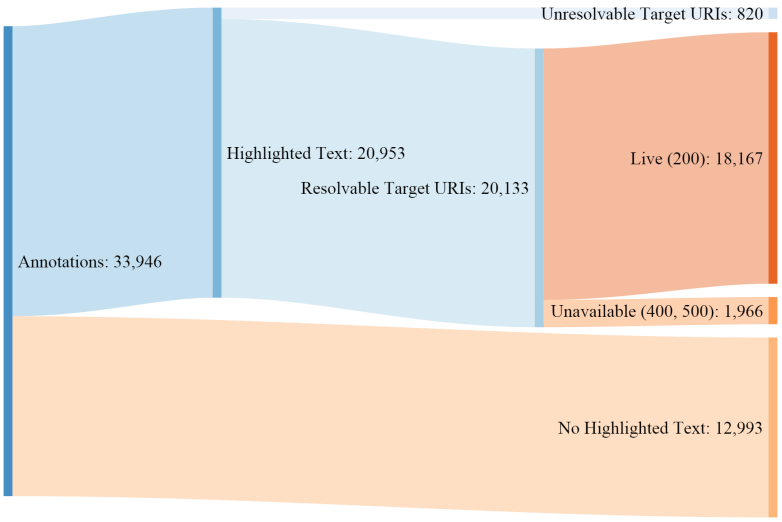

Fig. 9: HTTP Status Code of Annotations' Target URIs

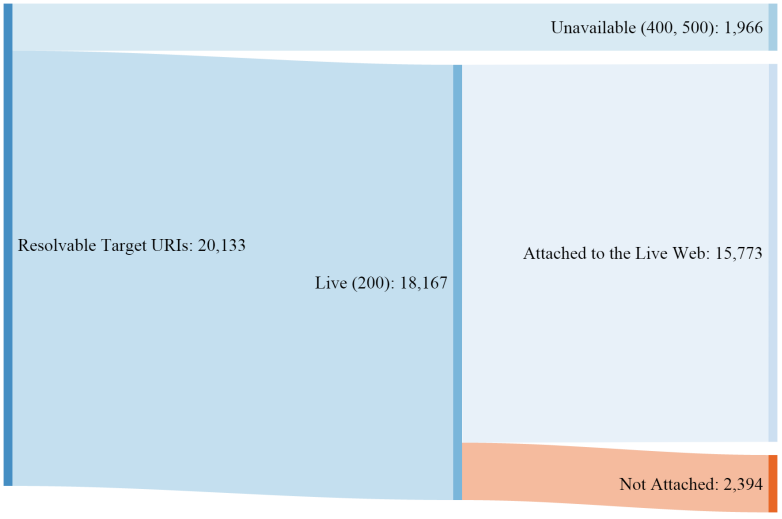

Fig. 10: Annotations that are Attached/Not Attached to the Live Web

88% of these annotations are orphans (19% of all annotations included in our analysis). As shown in Table 7, the majority of annotations have no mementos available at all. Those that can no longer be attached to their live web version are lost (orphans), but those that are still attached are in danger of being orphaned (41% of all annotations). These annotations can be recovered if these pages are archived before the annotated text changes. As we can see, 60% of annotations

Number of Annotations	Attached to Live Web Page	Attached to Memento (L)	Attached to Memento (R)
4091	Yes	Yes	Yes
93	Yes	Yes	No
100	Yes	No	Yes
182	Yes	No	No
251	No	Yes	Yes
69	No	Yes	No
44	No	No	Yes
156	No	No	No

Table 4: Annotation Targets with Existing Mementos Before and After the Annotation Creation Date.

are either orphaned or in danger of being orphaned. Figure 11 shows the status of current Hypothes.is annotations.

Table 8 shows the number of annotations that can be recovered using various archives, split by whether or not they are still attached to the live web. As expected the Internet Archive can be used to recover the most annotations.

Number of Annotations	Attached to Live Web Page	Attached to Memento (L)
1984	Yes	Yes
235	Yes	No
133	No	Yes
125	No	No

Table 5: Annotation Targets with Existing Mementos Only Before the Annotation Creation Date

5 Conclusions

In this paper, we analyzed the attachment of highlighted text annotations in Hypothes.is. We studied the prevalence of orphaned annotations, and found that 19% (3813) of the highlighted text annotations are orphans while 41% (8357) are in danger of being orphaned. We used Memento to look for archived versions of the annotated pages and found that 3% (547) of annotations that are not attached to the live web can be reattached to archived versions. We also found that for the majority of the annotations (11,273), no memento exists in the archives. This points to the need for archiving web pages at the time of annotation.

Number of Annotations	Attached to Live Web Page	Attached to Mementos (R)
1397	Yes	Yes
101	Yes	No
50	No	Yes
98	No	No

Table 6: Annotation Targets with Existing Mementos Only After the Annotation Creation Date

Number of Annotations	Attached to Live Web
7839	Yes
3434	No

Table 7: Annotation Targets with No Existing Mementos

Highlighted Text Annotations With Resolvable Target URIs: 20,133
Attached to Live Web: 15,773
Not Attached to Live Web: 4,360
Safe: 7,416
In Danger: 8,357
Reattached using Archives: 547
Orphaned: 3,813

Fig. 11: The Status of Current Hypothes.is Annotations

References

1. Agosti, M., Ferro, N., Frommholz, I., Thiel, U.: Annotations in digital libraries and collaboratories–facets, models and usage. In: Research and Advanced Technology for Digital Libraries, pp. 244–255 (2004)
2. Ainsworth, S.G., Alsum, A., SalahEldeen, H., Weigle, M.C., Nelson, M.L.: How much of the web is archived? In: Proceedings of the 11th ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 133–136. ACM (2011)
3. Aturban, M., Nelson, M.L., Weigle, M.C.: Quantifying orphaned annotations in hypothes.is. Proceedings of Theory and Practice of Digital Libraries (TPDL) pp. 15–27 (2015)

Archive	Attached to Live Web	Not Attached to Live Web
Internet Archive	6997 (94.30%)	455 (83.10%)
archive.is	679 (9.15%)	39 (7.12%)
Archive-It	562 (7.57%)	47 (8.59%)
github.com	80 (1.07%)	21 (3.83%)
wayback.vefsafn.is	71 (0.95%)	53 (9.68%)
arxiv.org	18 (0.24%)	0
www.webarchive.org.uk	4 (0.05%)	0
webarchive.loc.gov	3 (0.04%)	0
webarchive.nationalarchives.gov.uk	2 (0.02%)	0
discordia.wikia.com	1 (0.01%)	0
Total	8417 (113.40%)	615 (112.32%)

Table 8: Annotation Targets Recovered by Different Archives

4. Bar-Yossef, Z., Broder, A.Z., Kumar, R., Tomkins, A.: Sic transit gloria telae: Towards an understanding of the web's decay. In: WWW '04: Proceedings of the 13th International Conference on World Wide Web. pp. 328–337 (2004)
5. Brunelle, J.F., Kelly, M., Weigle, M.C., Nelson, M.L.: The impact of JavaScript on archivability. Proceedings of Theory and Practice of Digital Libraries (TPDL) pp. 1–23 (2013)
6. Brunelle, J.F., Weigle, M.C., Nelson, M.L.: Archiving deferred representations using a two-tiered crawling approach. In: iPRES 2015: Proceedings of the 12th International Conference on Digital Preservation
7. Brush, A., Bargeron, D., Gupta, A., Cadiz, J.J.: Robust annotation positioning in digital documents. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 285–292. ACM (2001)
8. Francisco-Revilla, L., Shipman, F., Furuta, R., Karadkar, U., Arora, A.: Managing change on the web. In: Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 67–76. ACM (2001)
9. Frommholz, I., Fuhr, N.: Probabilistic, object-oriented logics for annotation-based retrieval in digital libraries. In: Proceedings of the 6th ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 55–64. ACM (2006)
10. Furuta, R., Shipman III, F.M., Marshall, C.C., Brenner, D., Hsieh, H.w.: Hypertext paths and the world-wide web: Experiences with Walden's paths. In: Proceedings of the 8th ACM Conference on Hypertext. pp. 167–176. ACM (1997)
11. Haslhofer, B., Sanderson, R., Simon, R., Van de Sompel, H.: Open Annotations on multimedia web resources. Multimedia Tools and Applications 70(2), 847–867 (2014)
12. Haslhofer, B., Simon, R., Sanderson, R., Van de Sompel, H.: The Open Annotation Collaboration (OAC) model. In: Proceedings of the IEEE Workshop on Multimedia on the Web (MMWeb). pp. 5–9. IEEE (2011)
13. Klein, M., Van de Sompel, H., Sanderson, R., Shankar, H., Balakireva, L., Zhou, K., Tobin, R.: Scholarly context not found: One in five articles suffers from reference rot. PloS one 9(12), e115253 (2014)

14. Marshall, C.C.: Toward an ecology of hypertext annotation. In: Proceedings of the 9th ACM Conference on Hypertext and Hypermedia: Links, Objects, Time and Space—Structure in Hypermedia Systems. pp. 40–49. ACM (1998)
15. SalahEldeen, H.M., Nelson, M.L.: Losing my revolution: How many resources shared on social media have been lost? In: Proceedings of Theory and Practice of Digital Libraries (TPDL), pp. 125–137 (2012)
16. SalahEldeen, H.M., Nelson, M.L.: Resurrecting my revolution: Using social link neighborhood in bringing context to the disappearing web. In: Proceedings of Theory and Practice of Digital Libraries (TPDL). pp. 333–345 (2013)
17. Sanderson, R., Albritton, B., Schwemmer, R., Van de Sompel, H.: SharedCanvas: A collaborative model for medieval manuscript layout dissemination. In: Proceedings of the 11th Annual International ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 175–184. ACM (2011)
18. Sanderson, R., Ciccarese, P., Van de Sompel, H.: Designing the W3C open annotation data model. In: Proceedings of the 5th Annual ACM Web Science Conference. pp. 366–375. ACM (2013)
19. Sanderson, R., Van de Sompel, H.: Making web annotations persistent over time. In: Proceedings of the 10th ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 1–10. ACM (2010)
20. Soo, V.W., Lee, C.Y., Li, C.C., Chen, S.L., Chen, C.c.: Automated semantic annotation and retrieval based on sharable ontology and case-based learning techniques. In: Proceedings of the 2003 ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 61–72. IEEE (2003)
21. Van de Sompel, H., Nelson, M.L., Sanderson, R.: HTTP framework for time-based access to resource states – Memento, Internet RFC 7089. http://tools.ietf.org/html/rfc7089 (2013)
22. Van de Sompel, H., Nelson, M.L., Sanderson, R., Balakireva, L.L., Ainsworth, S., Shankar, H.: Memento: Time Travel for the Web. Tech. Rep. arXiv:0911.1112 (2009)